Innovazione via progetti

Daniele Spiga, Diego Ciangottini, Mirco Tracolli, Sara Vallero, Andrea Ceccanti, Daniele Cesini, Alessandro Costantini, Cristina Duma, Giacinto Donvito, Marica Antonacci, Stefano Nicotri, Davide Salomoni, Luciano Gaido

Outline

- General Intro
- INDIGO-IAM
- DODAS EOSC-HUB Thematic Service
- XDC
- DEEP
- A couple of connecting the dots examples
  - Continuous training in ML applications
  - The IoTwins project

General Intro

EOSC Ecosystem…the oversimplified story

EOSC-hub mobilizes providers from European major digital infrastructures, EGI, EUDAT CDI and INDIGO-DataCloud jointly offering services, software and data for advanced data-driven research and innovation

- 100 Partners, 76 beneficiaries (75 funded)
- €33M total budget
- 36 months: Jan 2018 – Dec 2020

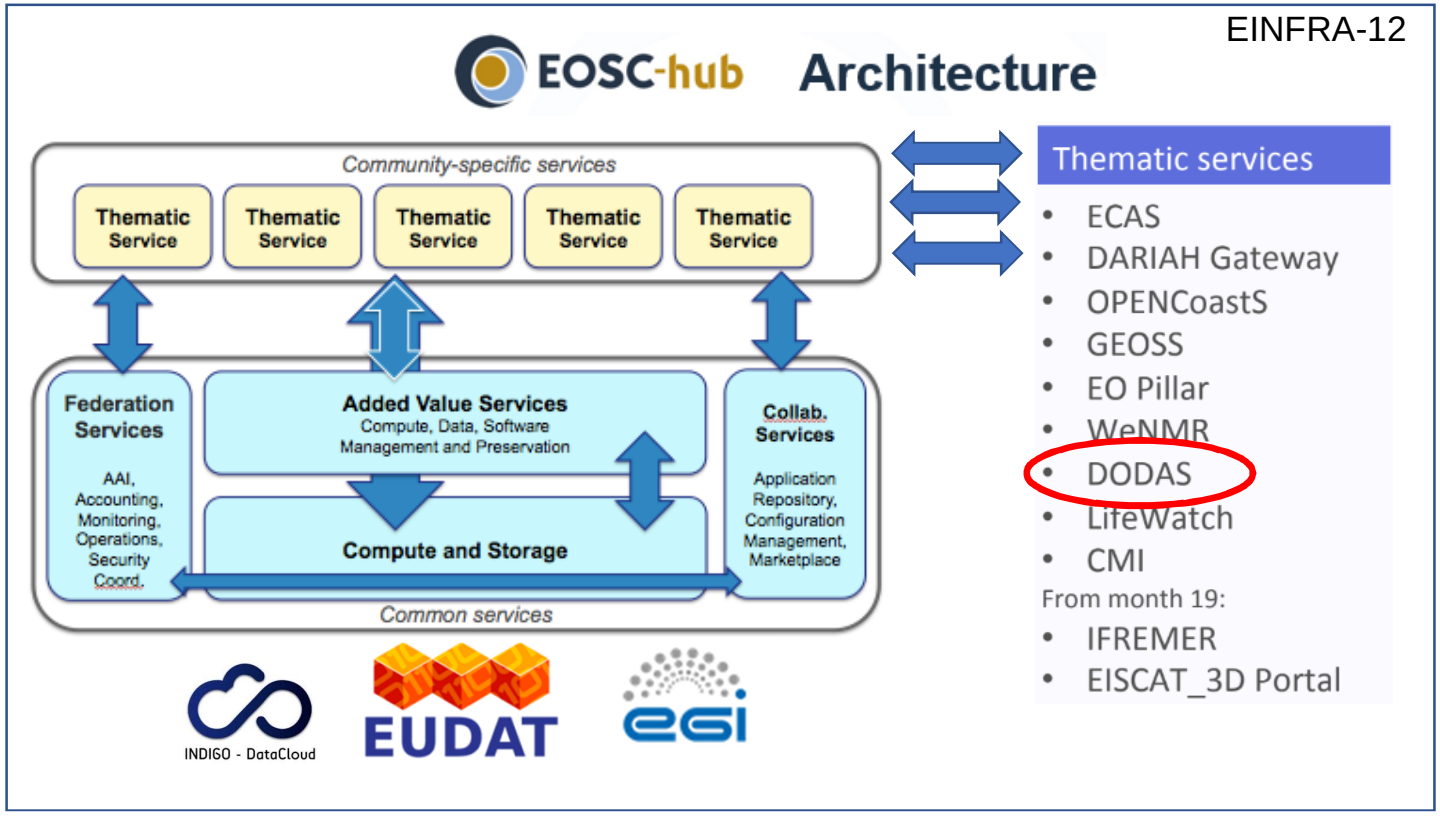

EOSC Ecosystem…the oversimplified story

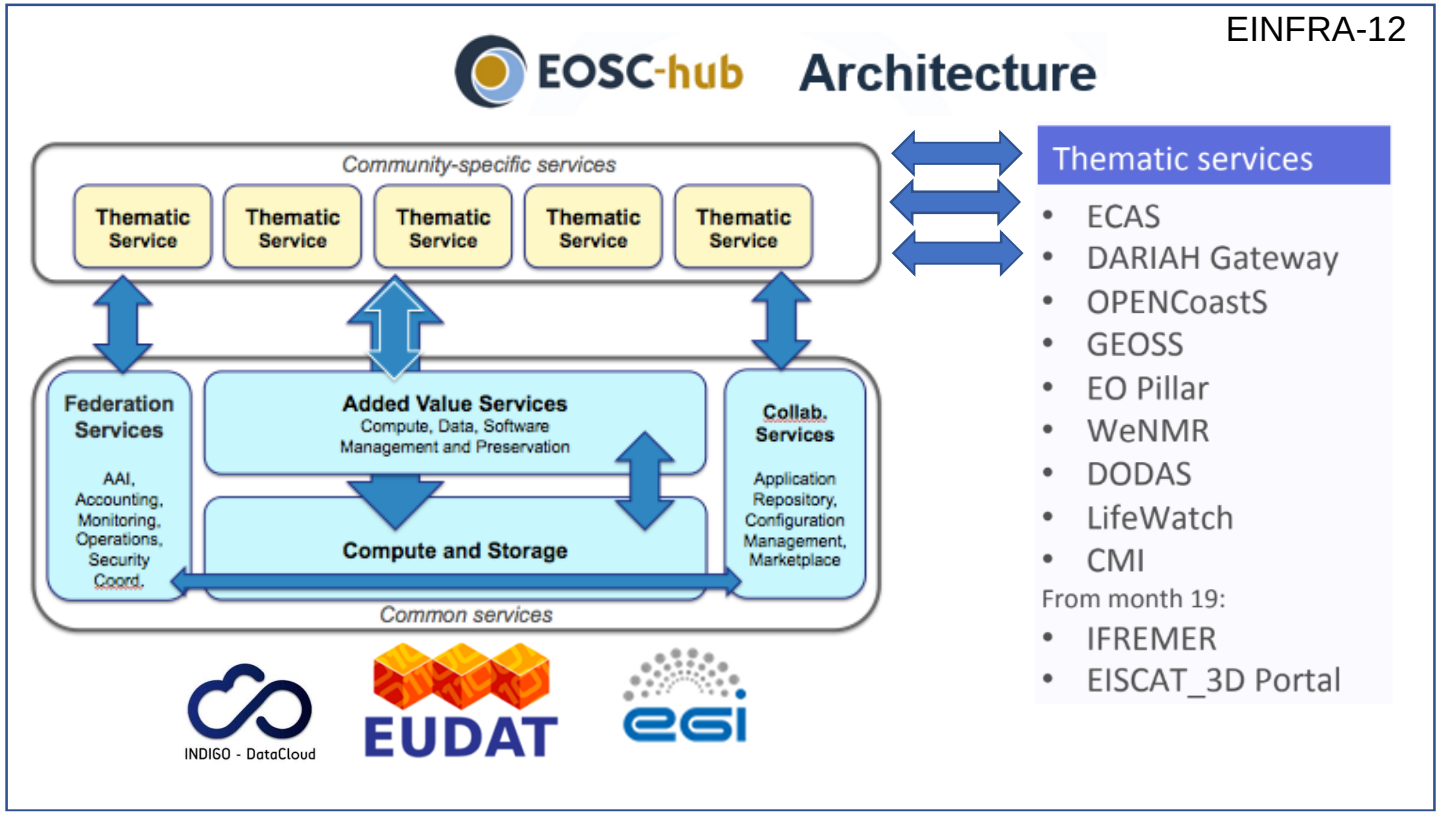

EOSC Ecosystem…the oversimplified story

EINFRA-21

Hybrid DataCloud

eXtreme DataCloud

DARE

EUXDAT

PROCESS

FREYA

EINFRA-12

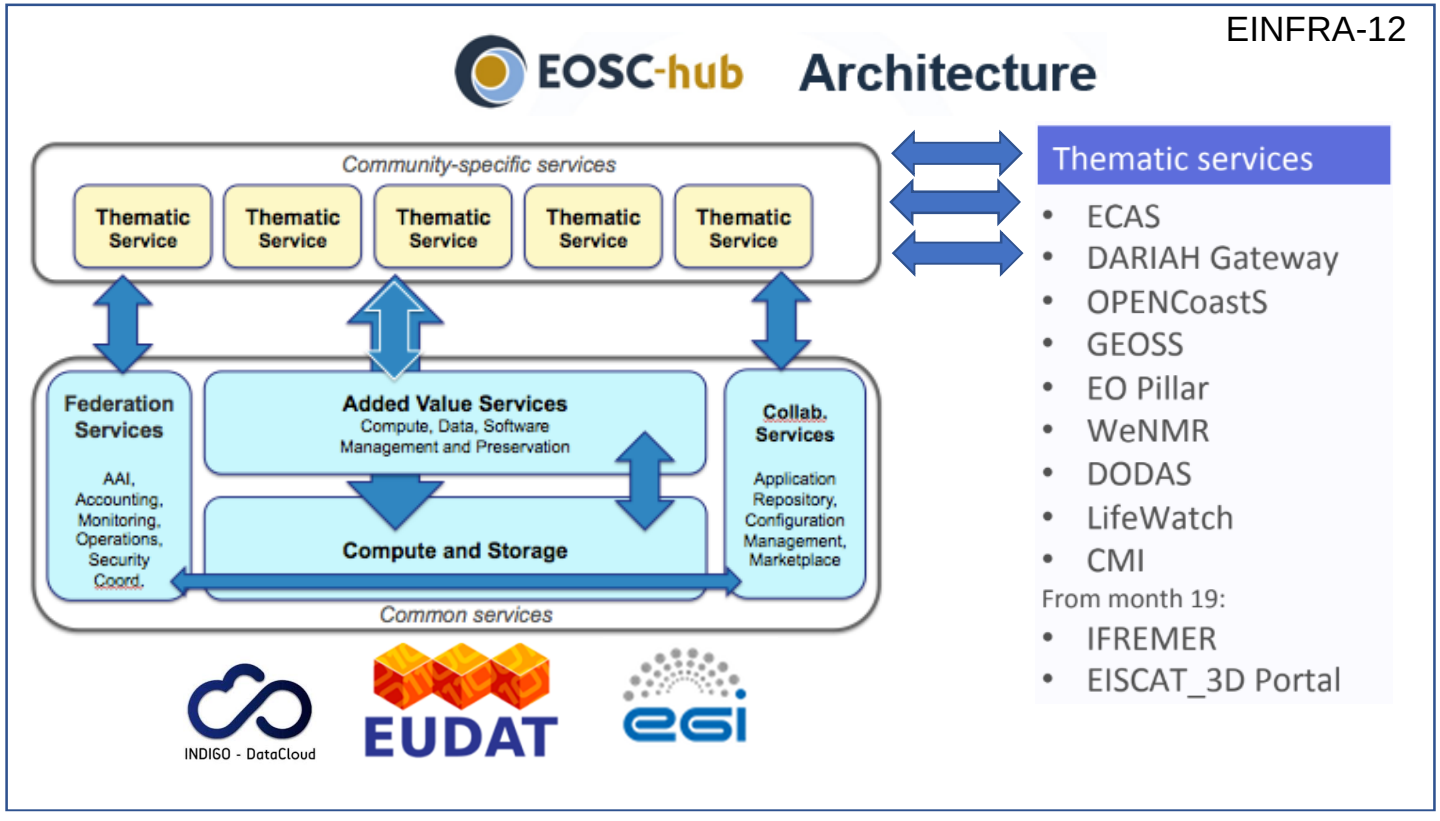

EOSC Ecosystem…the oversimplified story

EINFRA-21

Hybrid DataCloud

eXtreme DataCloud

DARE

EUXDAT

PROCESS

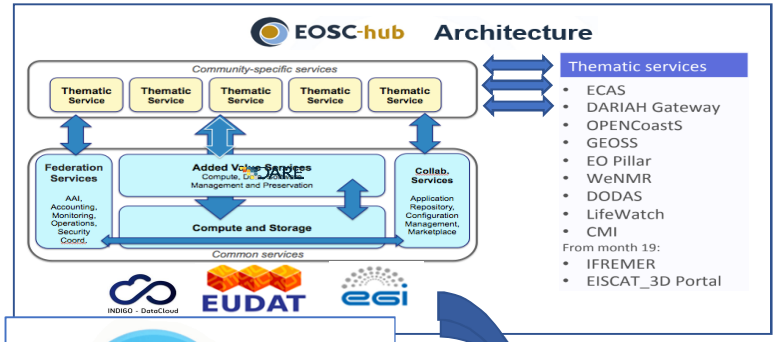

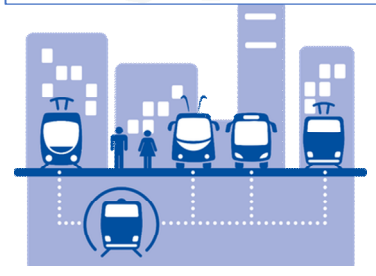

Communities/Experiments
E-infrastructures
Can or cannot be part of the EOSC
i.e. ESCAPE, HNSciCloud

Other infrastructures, not directly linked with the EOSC e-infra
i.e. SmartCities, IoT/bigData infra

ESCAPE

HELIX NEBULA THE SCIENCE CLOUD

INDIGO-IAM

Indigo-IAM

- Self contained, comprehensive Identity and Access Management solution
  - Originally developed within INDIGO-DC
  - Currently sustained by INFN for the foreseeable future with support from:
    - EOSC-HUB
    - ESCAPE
- Flexible Authentication Support
  - X509, OpenIDConnect, user/passwd
- Easy integration with off-the-shelf components thanks to OIDC
- VOMS support integrated
- **Selected by WLCG to be at the core of the next-generation WLCG authorization service in support of the LHC computing**

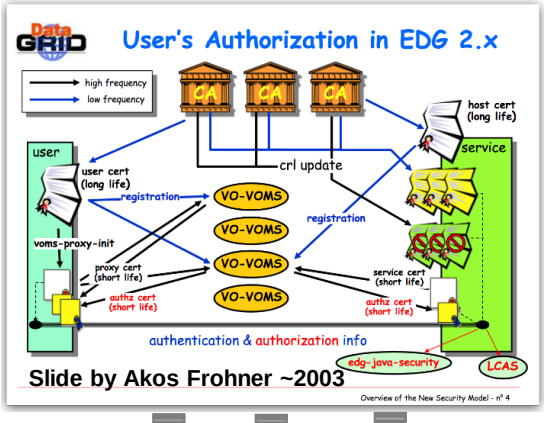

Slide by Akos Frohner ~2003

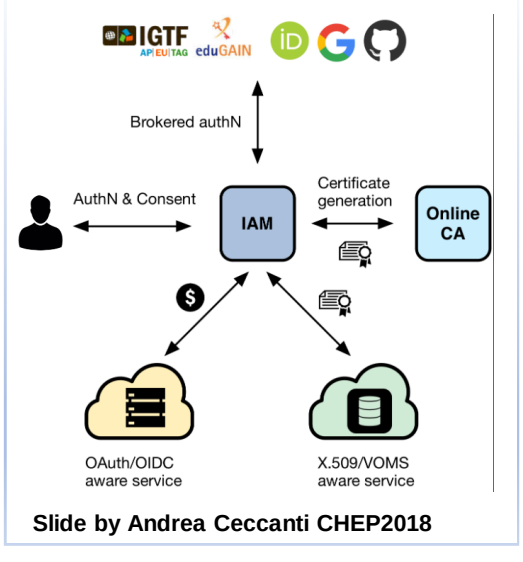

Slide by Andrea Ceccanti CHEP2018

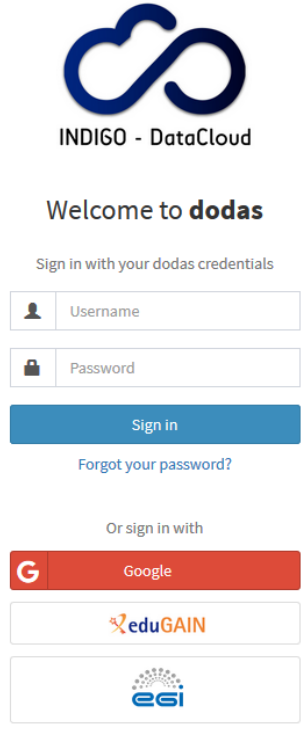

DODAS

- See also
  - DEMO: DODAS: HTCondor locale su INFN-CC + HTCondor distribuito (Daniele Spiga)
  - Talk on Friday 11:00 – Matteo Duranti:

    AMS and DAMPE: first experiences with federated cloud solutions and a look toward the future

DODAS in a nutshell

©D.Spiga

- DODAS: Dynamic On-Demand Analysis Service
  - A Thematic Service within EOSC-hub EU project service portfolio
- A open source deployment manager
- **Allows on-demand creation and configuration of container based clusters** for data processing with **almost zero effort**
- A cluster can be a standalone set of resources, a WLCG Tier*-like an extension of an existing center and more
  - BigData Analytics, Batch System as a Service, Distributed processing framework for ML
- **Support for hybrid clouds deployment**
- High level of automation and self-healing
- Supports communities-tailored (user-tailored) applications and software for data processing
- **Flexible Authentication and Authorization** model
- Based on “industry standards” to minimize code development and maintenance

DODAS Architecture

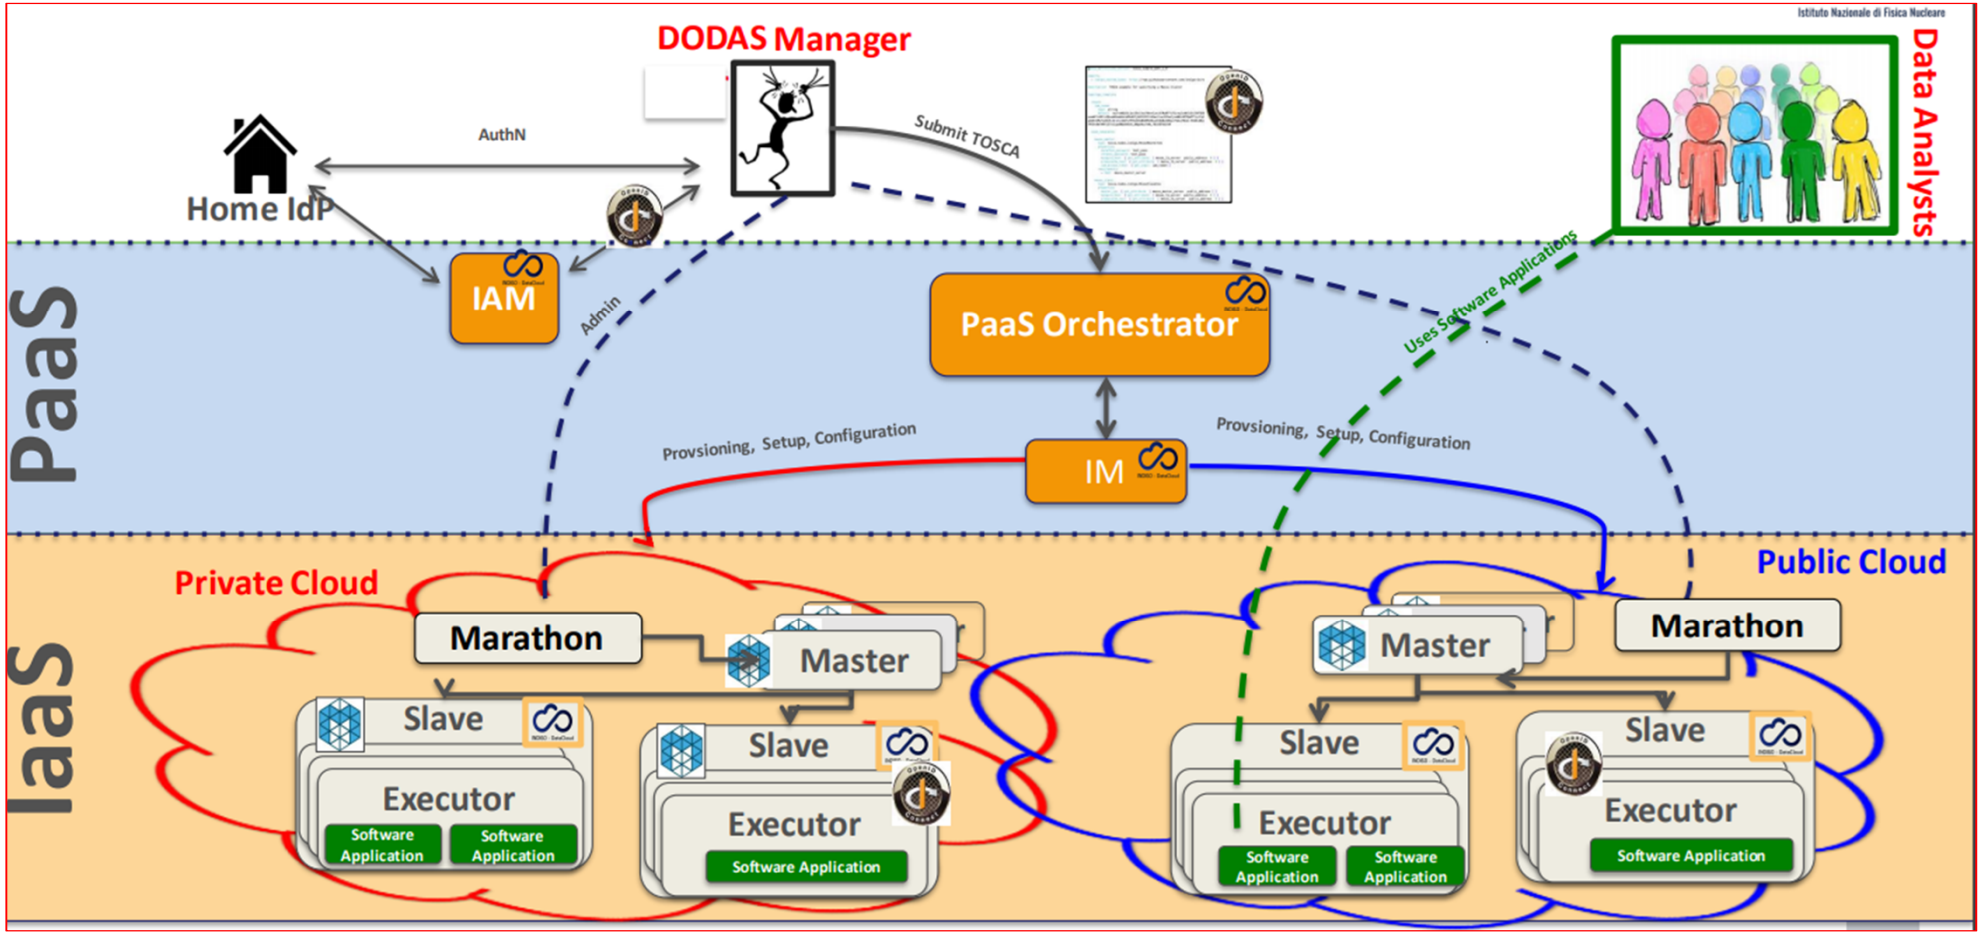

©D.Spiga

Implementing Vacuum with DODAS

- DODAS relies on vacuum approach to provide WLCG-like resources and it is responsible to automate:
  - Bare-hosts (e.g. VMs) instantiation based on user requirements, defined at TOSCA level
    - Virtual hardware can be scaled up/down (elasticity)
  - Services and software configurations at host leves (e.g. CVMFS, docker engine etc)
  - Container orchestrator deployment (e.g. K8s, Mesos/Marathon), and this is in form of dockers
  - Deployment and execution of services/microservices (e.g. Worker Nodes, squid proxy, x509 cache) over orchestrators
    - **this is how worker nodes are "spontaneously produced" and scaled up/down**
  - In addition DODAS provides a JWT based ecosystem for authentication and authorization: INDIGO-IAM

- Current incarnation is based on HTCondor as a mean to manage (aka overlay) distributed worker nodes (startd/glideins)
  - Does not deploy Computing Element (CE) but it could be added (example of modularity)
  - DODAS Vacuum system is integrated in the **CMS Computing infrastructure aka HTCondor Global pool** (see next slides)

DODAS is fully integrated into the CMS computing model to create lightweight ephemeral WLCG-Tier on demand

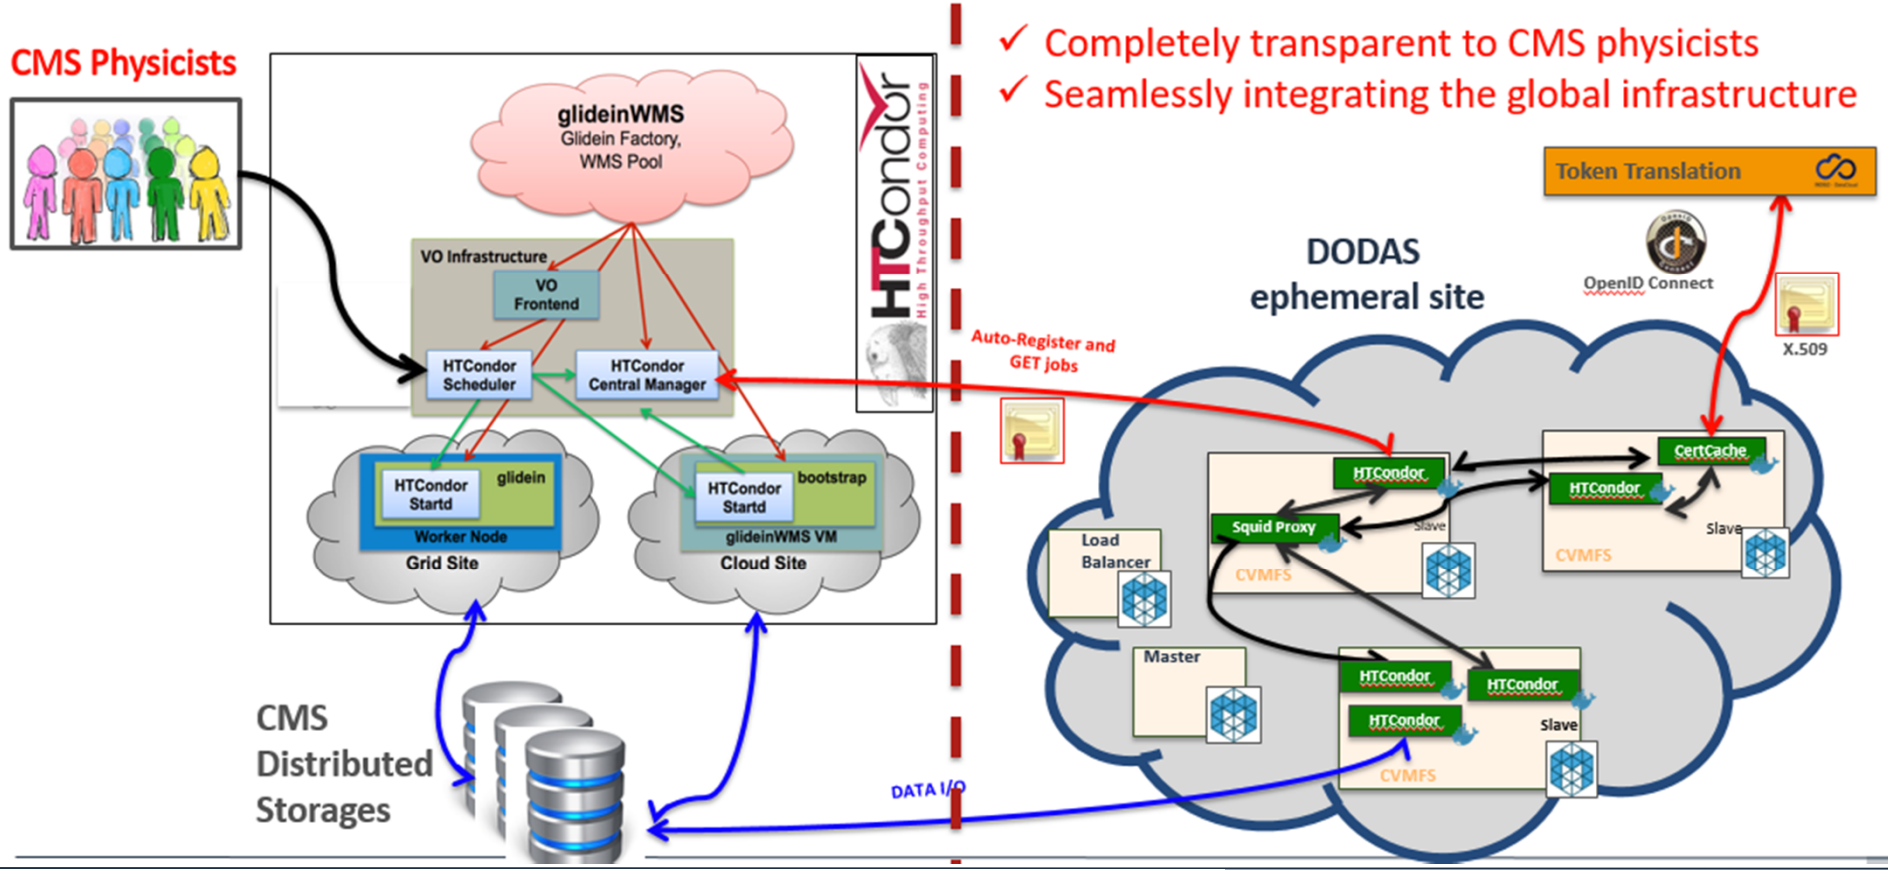

DODAS and public cloud: CMS analysis jobs

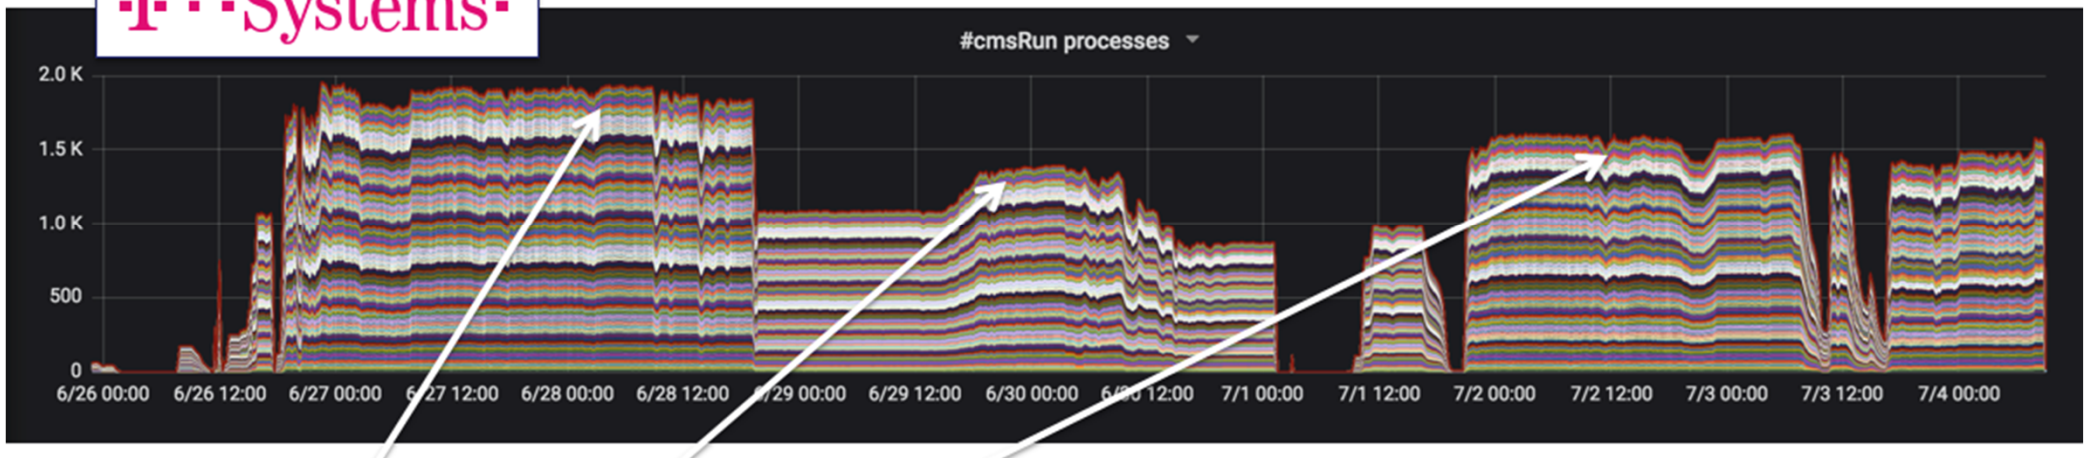

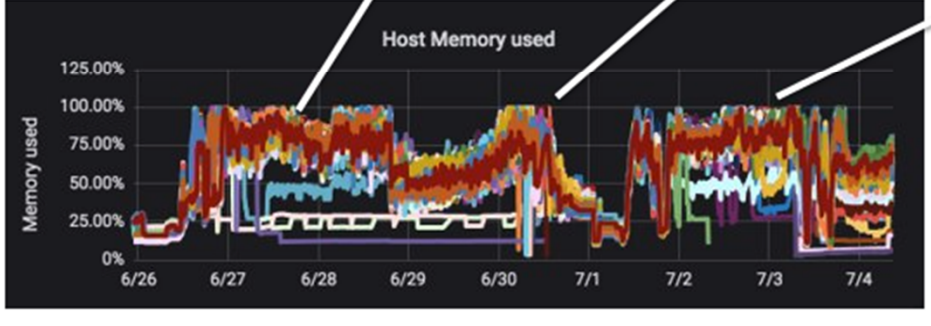

- Elasticity and self-healing
- Stability over days/weeks (120k jobs)
- Handling “special requirements” high memory jobs

DODAS in the Real Life @ CMS

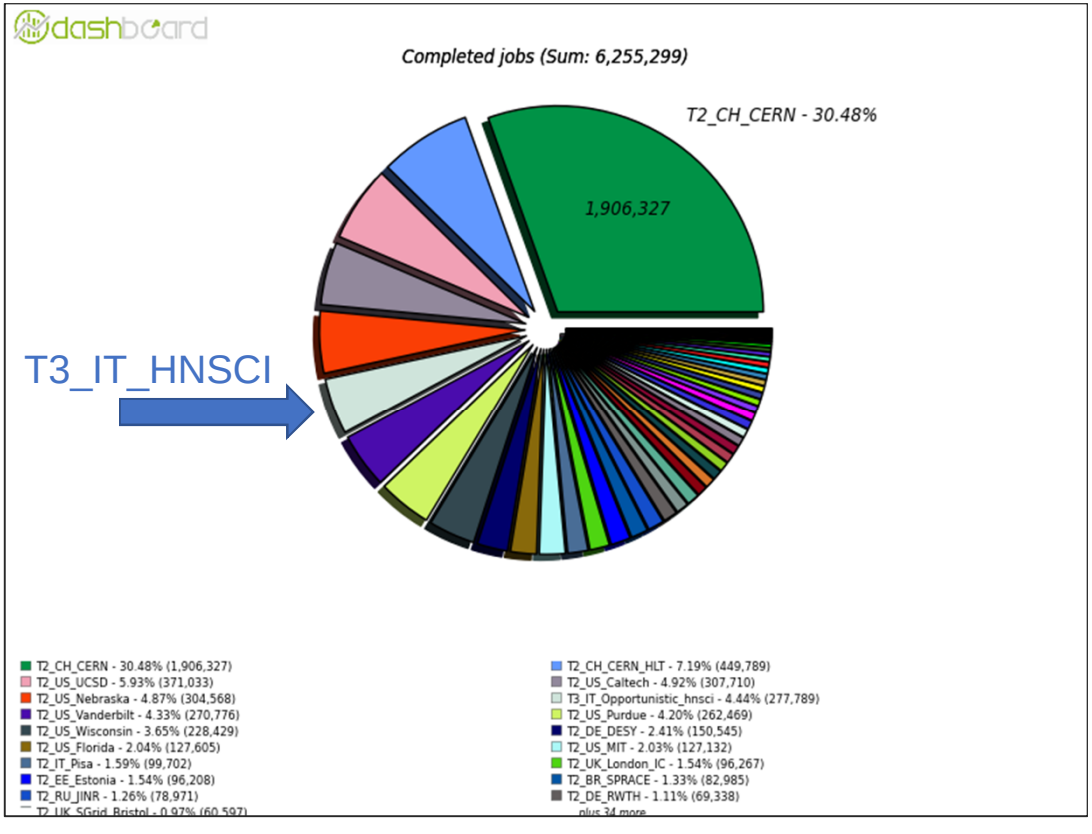

DODAS within top 6 CMS Tier2s for 10 days

CMS users analysis: hh—>2b2tau

→ Skimming minAOD (Data & MC) to produce flat-ntuple

- Francesco Brivio INFN
- Chiara Amendola LLR

DODAS @ VIRGO

- Configuration of a DODAS HTCondor cluster for Virgo workloads is being defined
- **What we have so far:**
  - Test cluster deployed at Recas Bari
  - Software packaged in Singularity images and distributed via CVMFS
  - *Dockerized* CVMFS server (stratum0/1) deployed in Torino
  - Workflows with auto-generated data being tested
- **What is coming:**
  - Data access through custom CVMFS acting as file catalog to remote data @ CNAF
  - Cache layer @ CNAF (see Virgo talk on Friday)
  - Scale out the HTCondor cluster on the Torino OpenNebula Cloud

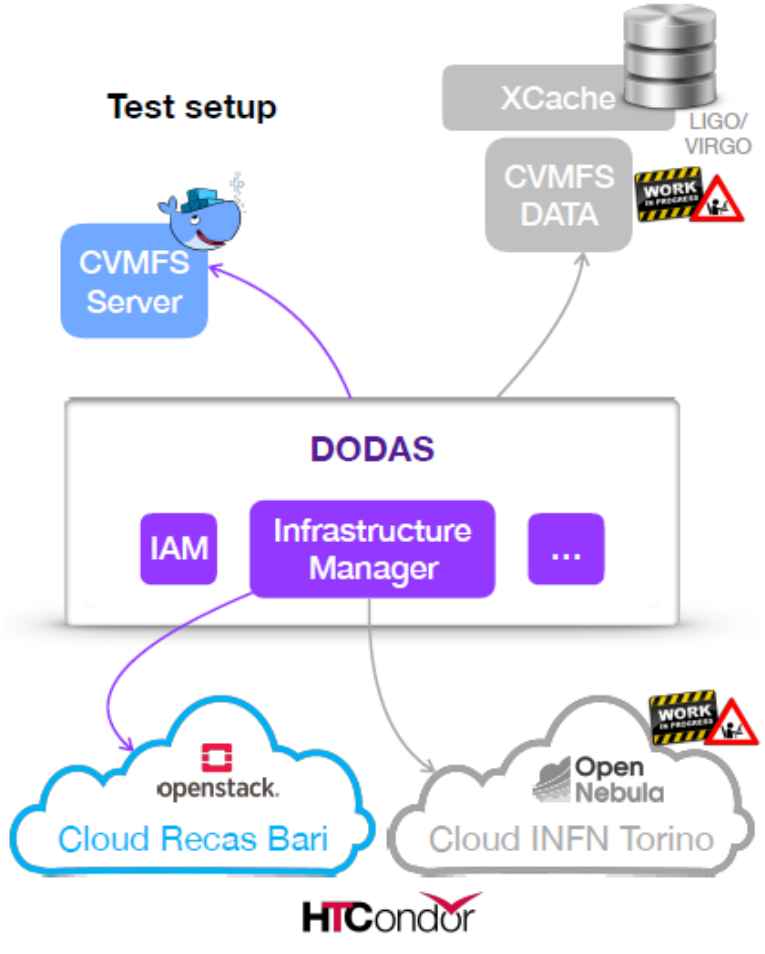

XDC – eXtreme-DataCloud

- See also
  - DEMO: Demo sulle dynamic-cache (Diego Ciangottini)
  - Talk on Friday 09:10 – Diego Ciangottini:
    Integration of a smart Italian cache federation for CMS

XDC

- The eXtreme DataCloud is a software development and integration project
- Develops scalable technologies for federating storage resources and managing data in highly distributed computing environments
  - Focus on efficient, policy driven and Quality of Service based DM
- The targeted platforms are the current and next generation e-Infrastructures deployed in Europe
  - European Open Science Cloud (EOSC)
  - The e-infrastructures used by the represented communities
- Addresses the EINFRA-21-2017 (b)-2: “Computing e-infrastructure with extreme large datasets”
  - Deal with heterogeneous datasets
  - Bring to TRL8 and include in a unified service catalogue services and prototype at least at TRL6
- 3.1M€ Nov2017 - Jan 2020, 8 partners, 7 countries

The Approach

- **Improve already existing, production quality Data Management services**
  - By adding missing functionalities requested by research communities
  - Based mainly on technologies provided by the partners and by the INDIGO-Datacloud project
  - Must be coherently harmonized in the European e-Infrastructures

INDIGO PaaS Orchestrator

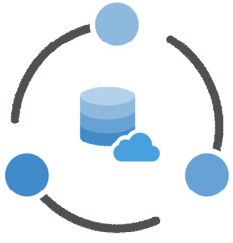

INDIGO CDMI Server

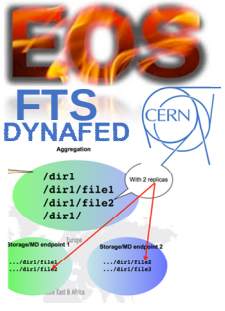

A User Driven Project

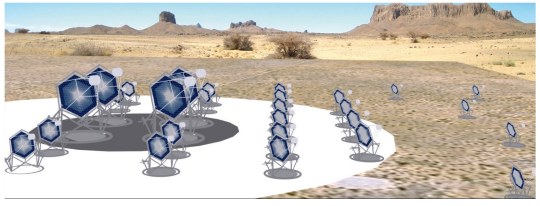

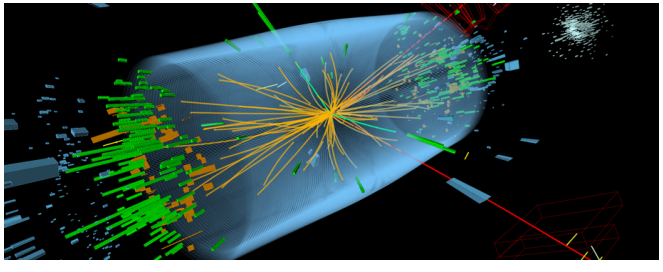

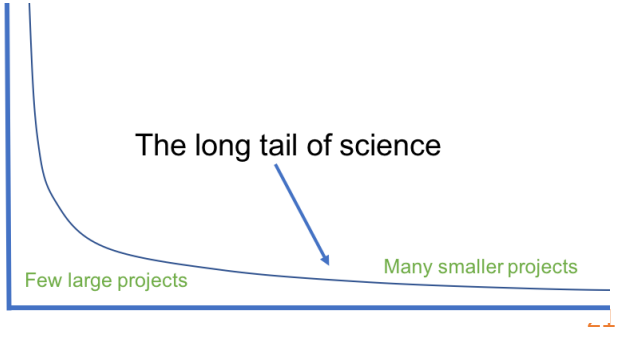

The New Functionalities

- Intelligent & Automated Dataset Distribution
  - Orchestration to realize a policy-driven data management
  - Data distribution policies based on Quality of Service (i.e. disks vs tape vs SSD) supporting geographical distributed resources (cross-sites)
  - Data lifecycle management
- Data pre-processing during ingestion
- Metadata management
- Data management based on storage events
- Smart caching
  - Transparent access to remote data without the need of a-priori copy
    - To support dynamic inclusion of diskless sites
    - To improve efficiency in multi-site storage systems and storage federations (i.e. Datalakes)
- Sensitive data handling
  - secure storage and encryption

The XDC General Architecture

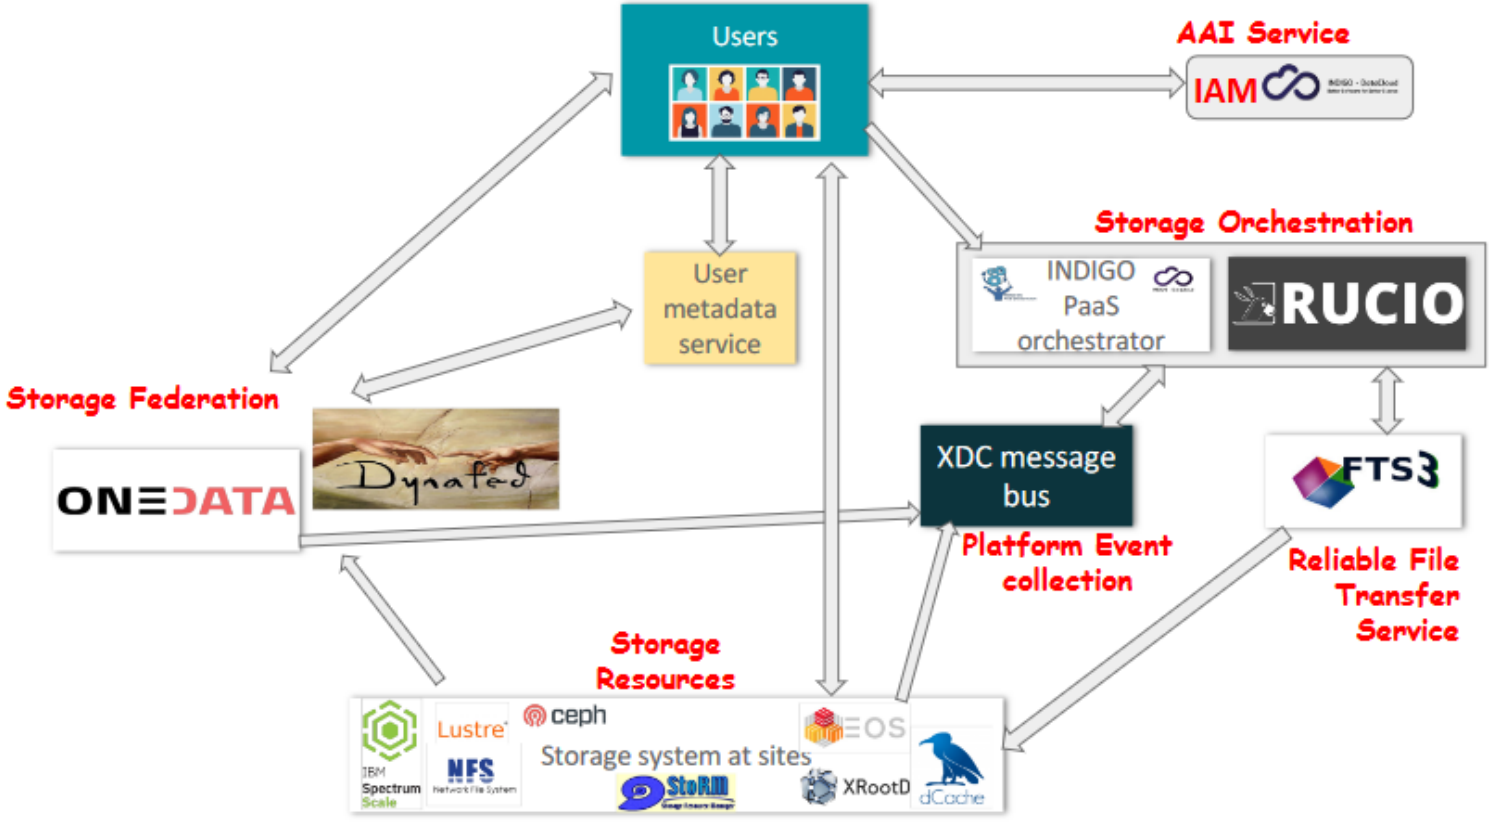

First XDC Release

- Involved tools
  - CachingOnDemand
  - dCache
  - Dynafed
  - EOS
  - FTS,GFAL
  - Onedata
  - PaaS Orchestrator plugin
  - TOSCA types & templates plugin
- Key technical highlights
  - OpenIDConnect support for token based authentication
  - New QoS types integration and support in dCache, FTS, GFAL
  - Orchestrator integration with other components
  - Performance improvements in Onedata
  - Support for groups and roles in Onedata
  - EOS-dCache integration
  - Caching systems instantiation
  - Storage events notification in dCache
  - EOS caching with XCache for geographic deployment
  - EOS external storage adoption

XDC-1/Pulsar

https://releases.extreme-datacloud.eu/en/latest/releases/pulsar/index.html

XDC Caching systems

Key technical highlights

- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- **Caching systems instantiation**
- Storage events notification in dCache
- EOS caching with XCache for geographic deployment
- EOS external storage adoption

- Deployment of Geo-distributed caches
- Network of unmanaged storage for hot data
- On-demand cache resources
- Based on xCache

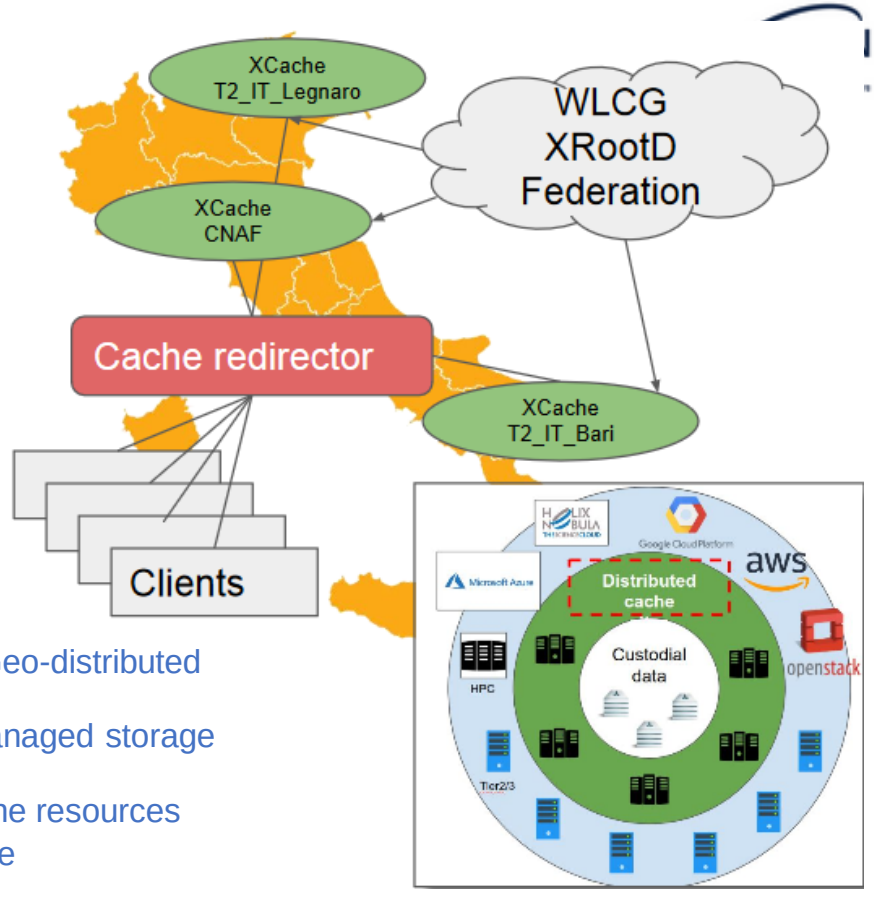

© Diego Ciangottini

CachingOnDemand provisioning

- Deploying methods for caching systems on **cloud/HPC resources:**
  - **ephemeral storage for hot data**
    - remote read latency
      - **Improve I/O inefficiency on WAN**
- **Available systems:**
- K8s, Docker, Ansible

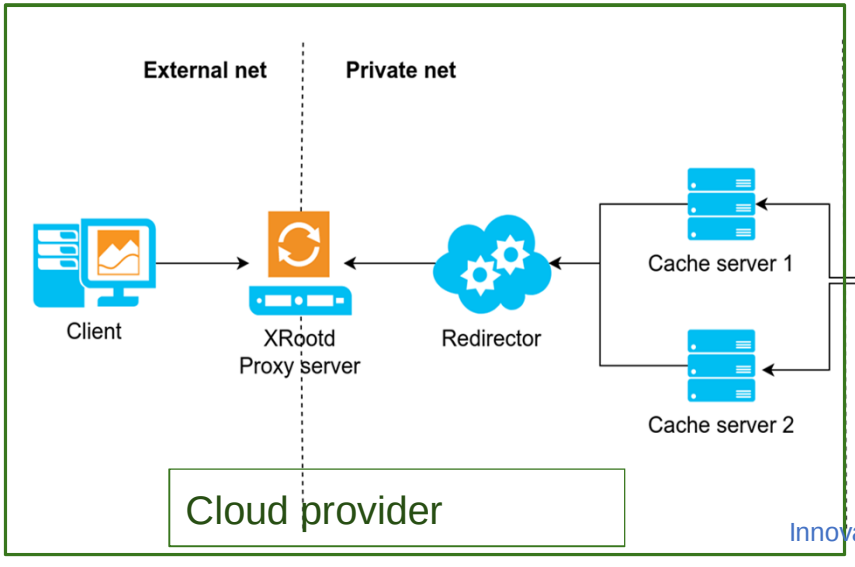

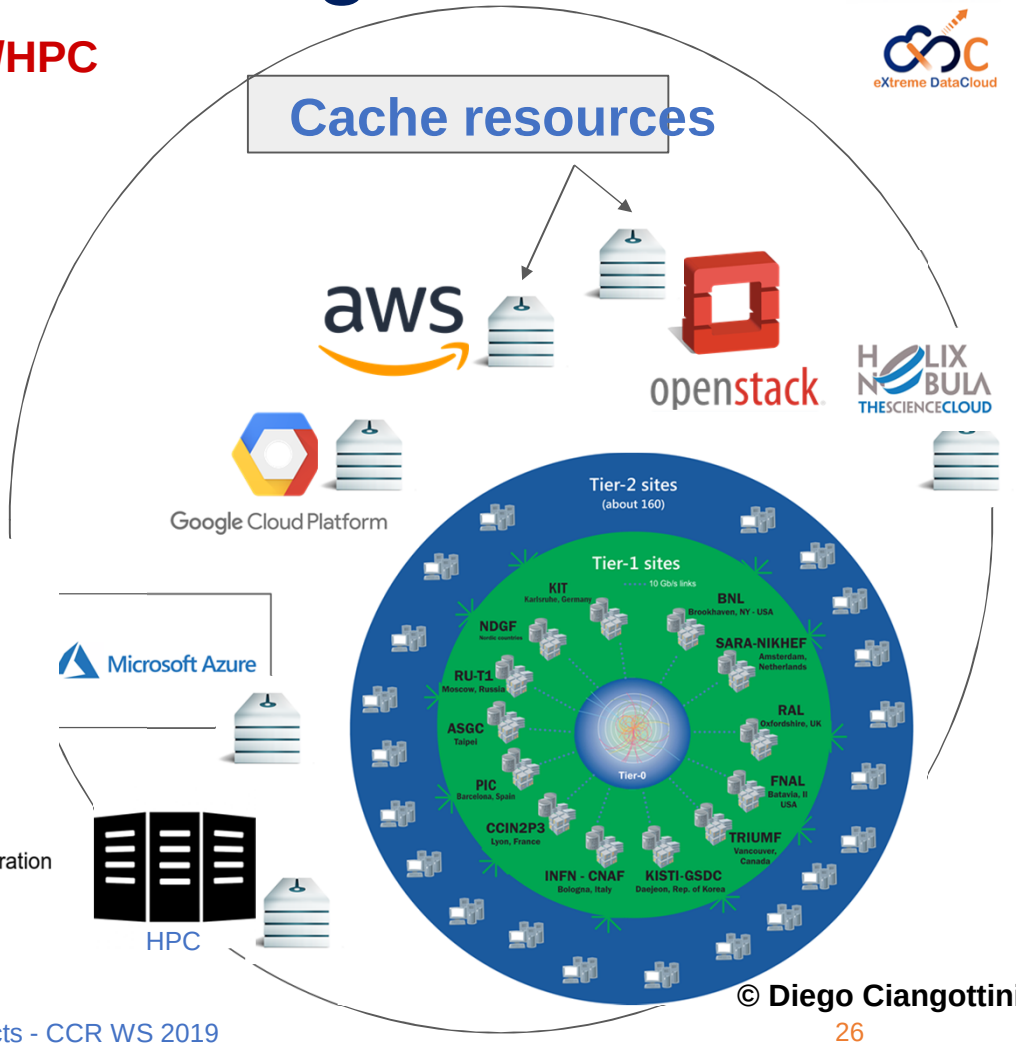

First XDC Release – Storage Notification

Key technical highlights

- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- **Storage events notification in dCache**
- EOS caching with XCache for geographic deployment
- EOS external storage adoption

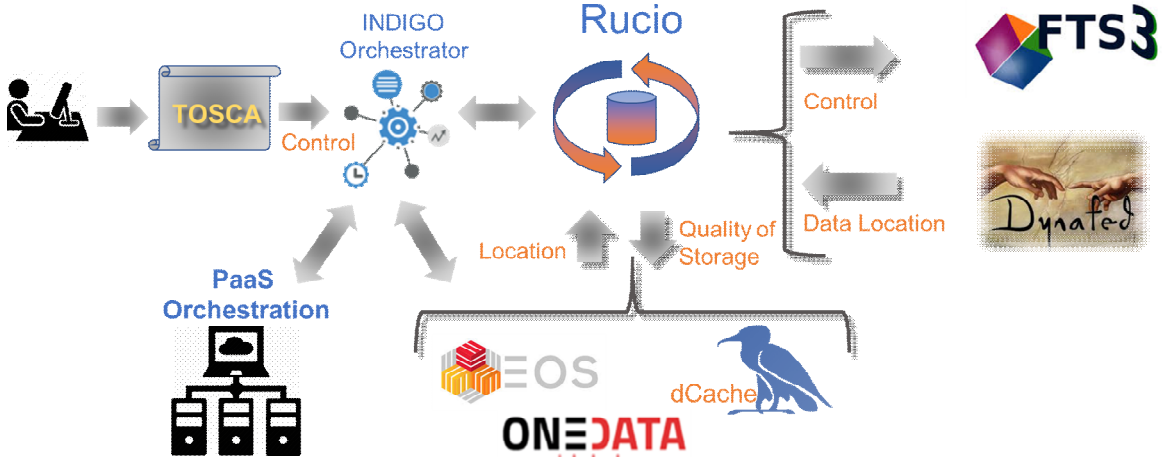

First XDC Release

Key technical highlights

- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- **Storage events notification in dCache**
- EOS caching with XCache for geographic deployment
- EOS external storage adoption

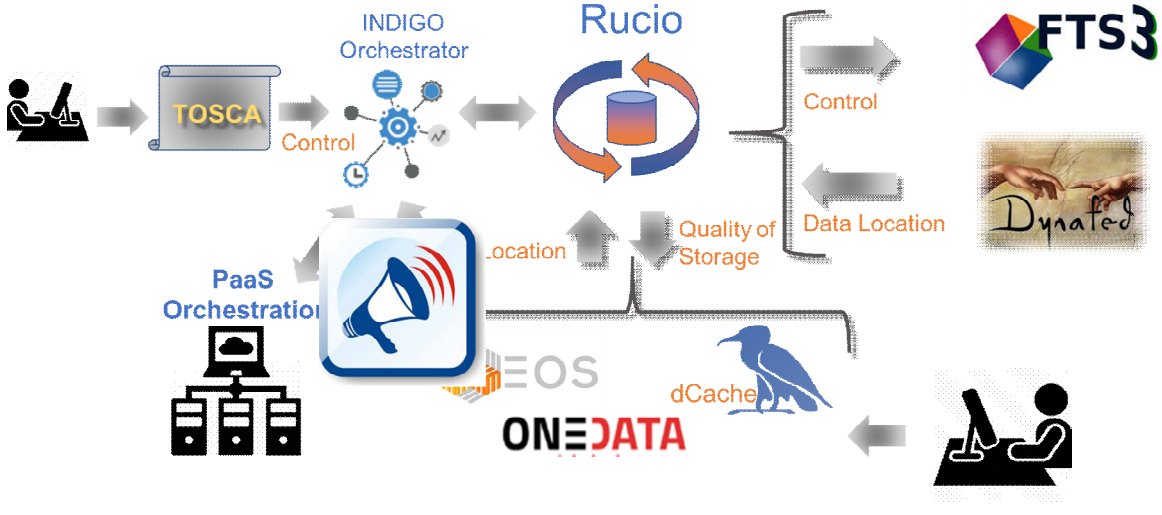

DEEP-HybridDataCloud

- See also
  - DEMO: Deployment di modelli per Machine Learning con utilizzo di GPU (Stefano Nicotri, Marica Antonacci)

DEEP serves different users

Basic Users: want to use a trained deep neural network for prediction/classification of their own data
→ no expert knowledge of ML
→ no access to high-level computing resources

docker

DEEPaaS API

Intermediate Users: want to use a trained deep neural network and adapt it for solving their problem (transfer learning)
→ some knowledge of ML
→ access to limited computing resources

docker Nextcloud

DEEPaaS API

Advanced Users: want to develop their own deep neural network (with special requirements as e.g. data privacy)
→ expert knowledge of ML
→ need access to high-level computing resources

docker Nextcloud COOKIECUTTER

DEEPaaS API

Selection of DEEP Use Cases

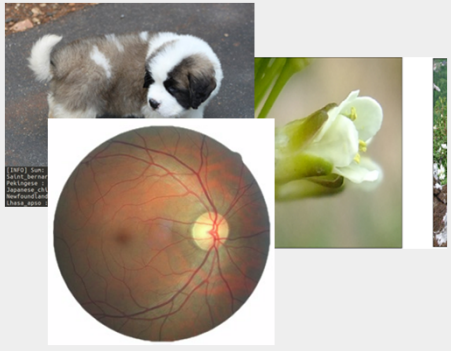

Deep Learning

- classical use case
- needs CPU/GPU resources
- various scenarios
  - *- web services*
  - *- transfer learning*
  - *- data privacy*

Post-processing of large datasets

- extremely large datasets
- executed in HPC environment

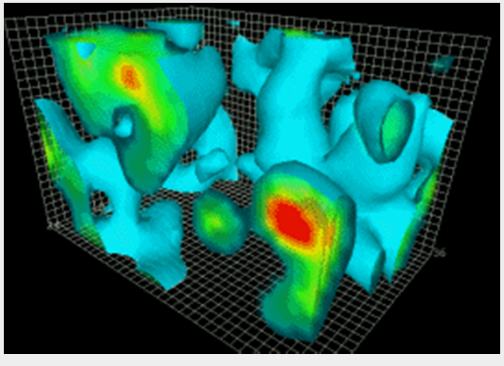

www.physics.adelaide.edu.au

Machine Learning on data streams

- hard time constraints
- execution within protected environments

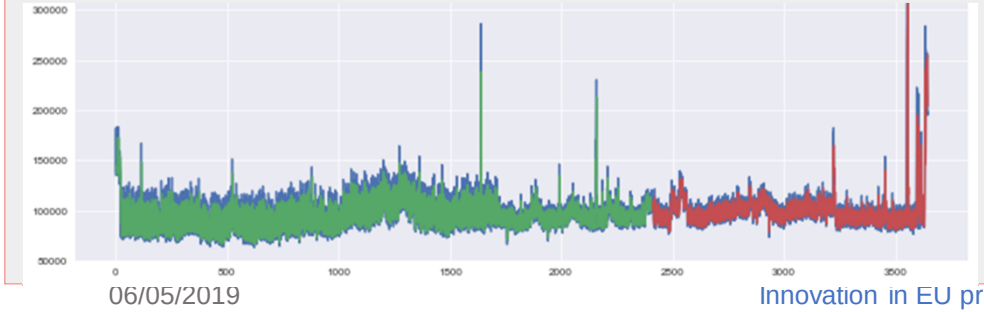

Use Case Plants



task: classification of plants

data: PlantNet

model: Xception

framework: Keras on TensorFlow


Vlaams Instituut voor de Zee

mncn museonacionaldecienciasnaturales

adopted by external partners

Use Case Retinopathy

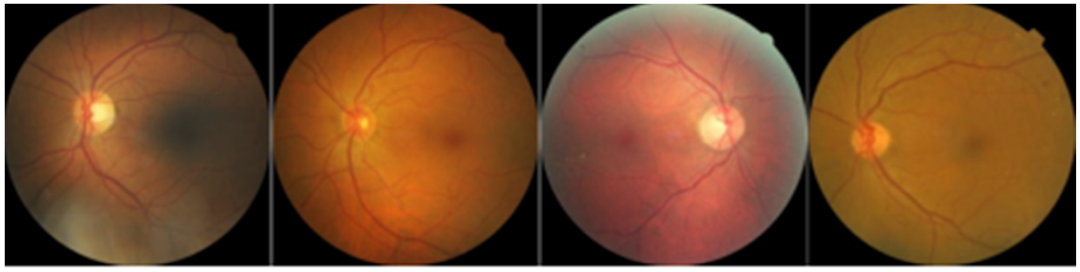

task: classification disease state

data: EyePACS

model: ResNet50

framework: TensorFlow

Use Case has been presented as Kaggle challenge (*benchmark dataset*)

application serves as prototype for data privacy restrictions
(→ distributed learning)

Use Case MODS

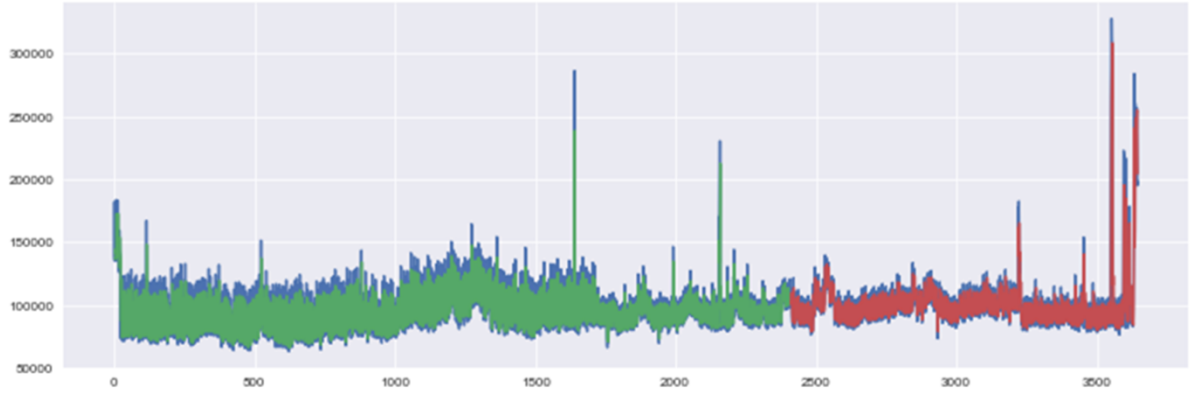

task: intrusion detection

data: EyePACS

model: LSTM/GRU

framework: Keras on TensorFlow

implementation of
- *Recurrent Neural Networks* variants,
- *Convolutional Neural Networks* and
- Fully Connected model (the latter two for comparison)

real-world scenario will provide severe time constraints

data privacy constraints still need to be addressed (data preprocessing)

DEEP Architecture

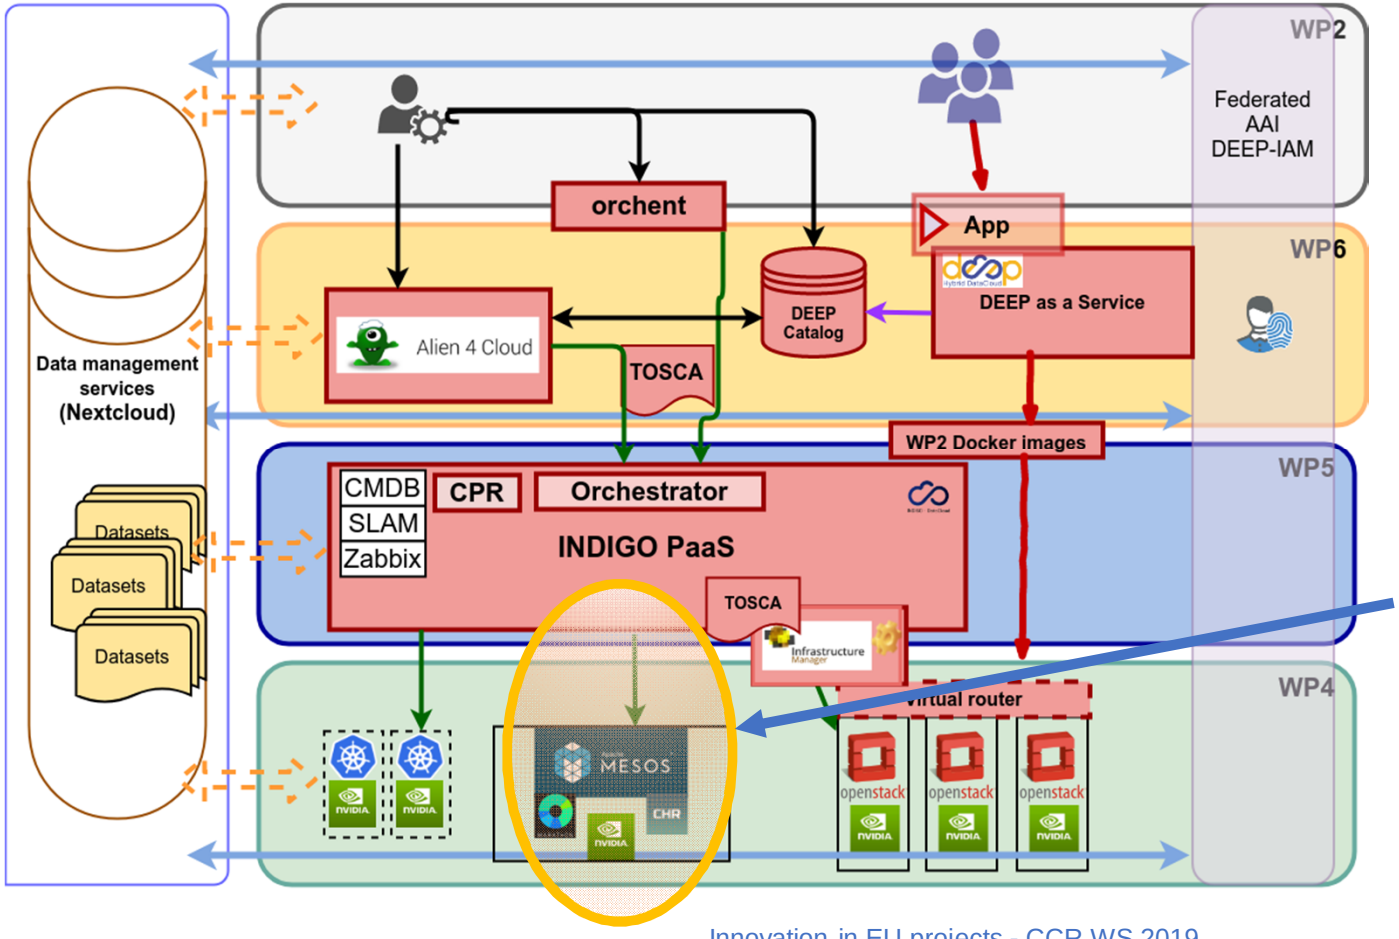

This will be demonstrated live soon after this talk!

- DEEPaaS to provide ML framework «as a service»
- Apache Mesos and Marathon used to orchestrate long-running services on containers
- Docker and Mesos UCR containers
- Instantiation on GPU resources

DEEP Orchestrator Dashboard

key features

- abstracts from technical details

- fully automatic

- to be used with DEEP user account only

Alien4Cloud Composition

Automated composition of the TOSCA Template from the GUI

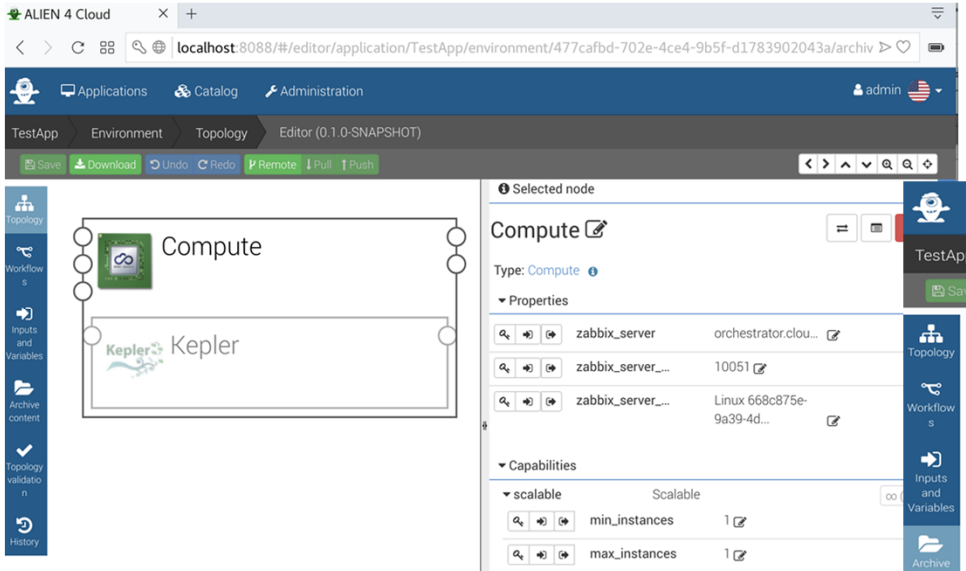

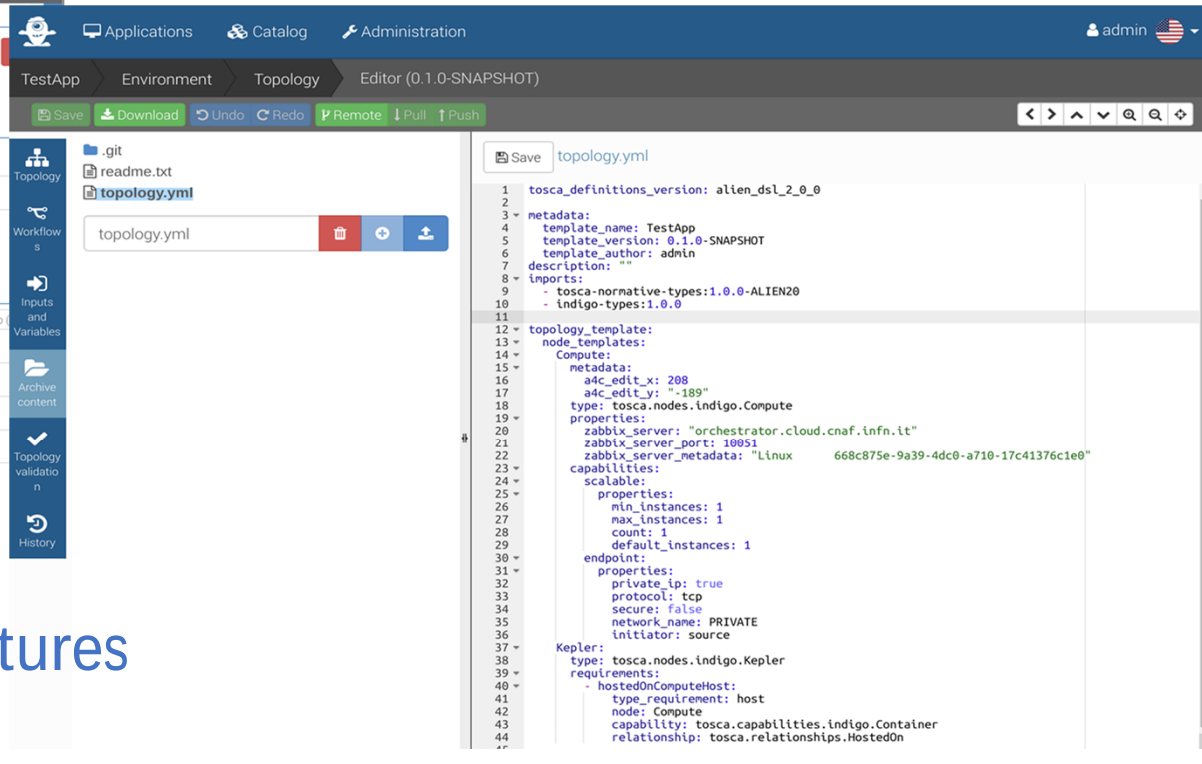

From simple to complex architectures

https://www.youtube.com/playlist?list=PLJ9x9Zk1O-J_UZfNO2uWp2pFMmbwLvzXa&jct=Ukj6uXSUFDnebigzIqf-0SsQhxiliQ&disable_polymer=true

Main functionalities released in the DEEP-1

- Increasing the **TRL up to 8** for the functionalities/services released
- A trial-and-error mechanism to **re-schedule the deployment** on the next available cloud provider from the list of candidate sites
- The capability to provision and **configure resources on** networks that only provide **private IP** addressing
  - only the front-end requires a public IP, to be accessed by the user, while the working nodes can be assigned a private IP address
  - Public IPs are typically a scarce resource
  - reducing the level of exposure of working nodes to the Internet reduces the chance of external intrusions
- Support for **GPU and Infiniband** resources
- Hybrid deployments on **multiple sites**
- Support for specifying **specialized computing hardware** improved support for deployment failures

Continuous training for ML applications

- See also:
  - DEMO: Spark/Hadoop per inferenza (Mirco Tracolli)

ML workflow for continuous training

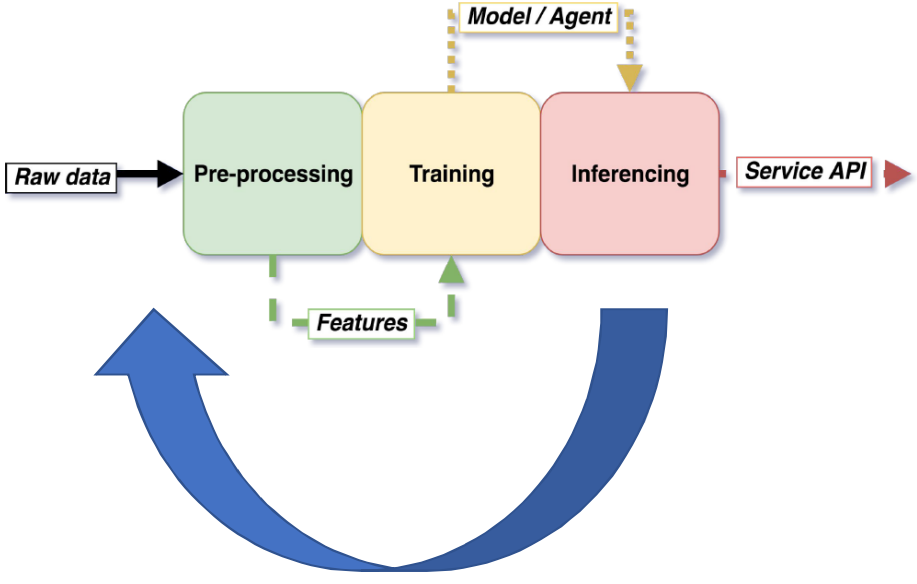

DEMO: Spark/Hadoop per inferenza (Mirco Tracolli)

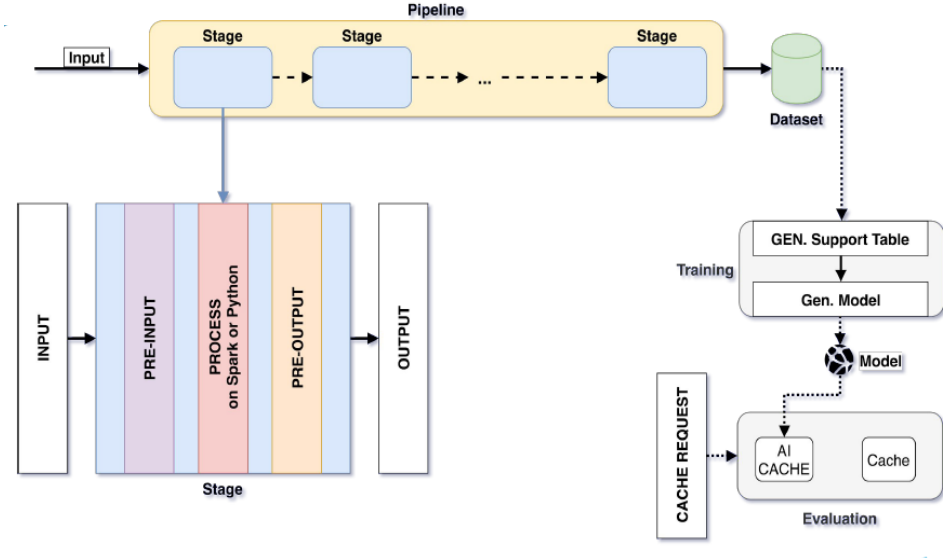

Continuos ML workflow with XDC, DEEP and DODAS

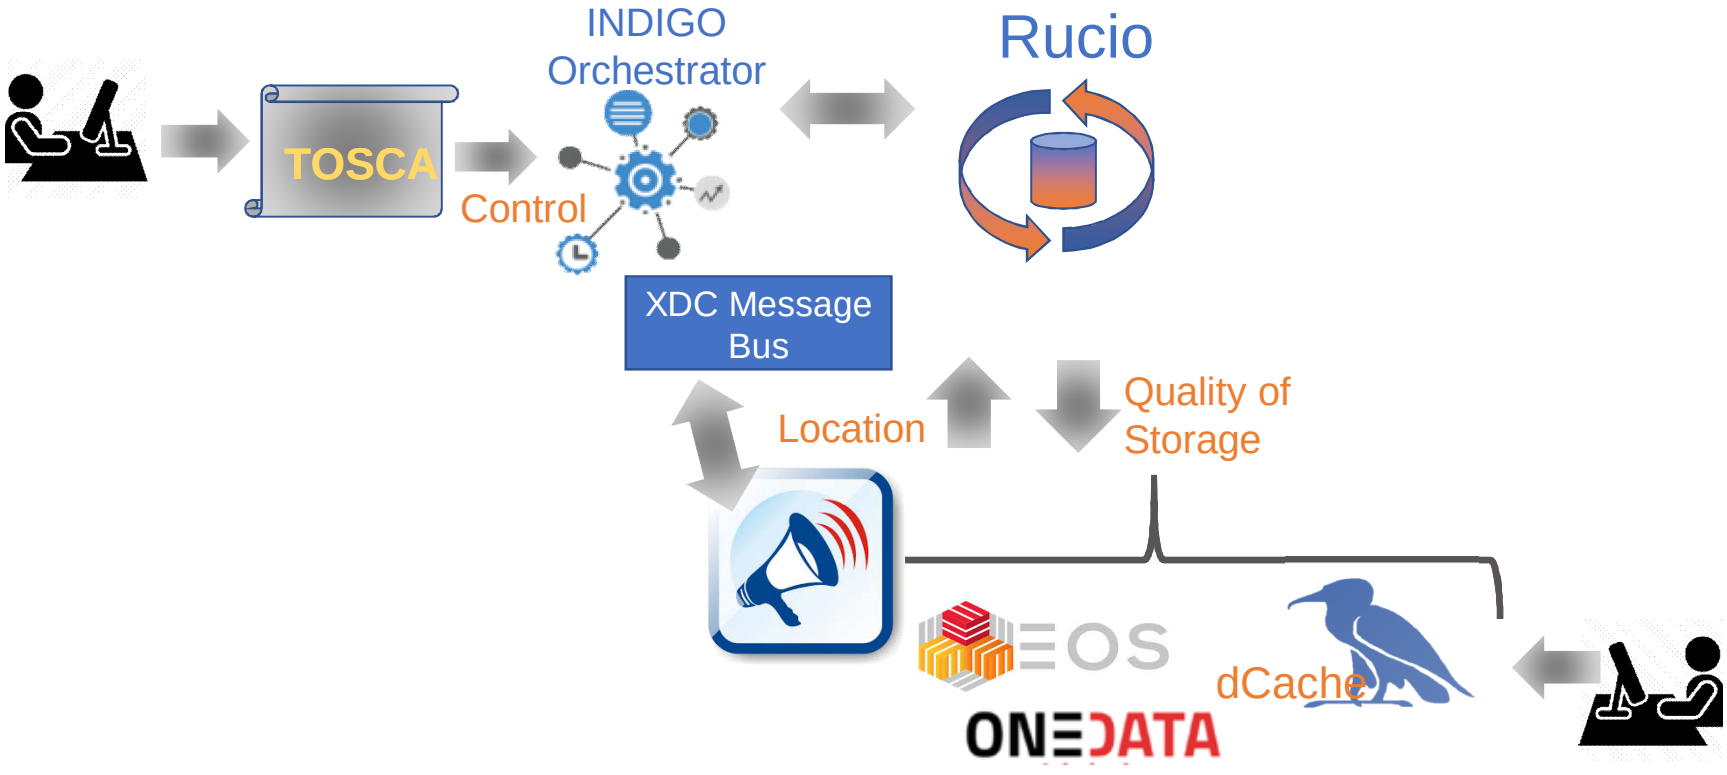

Continuos ML workflow with XDC, DEEP and DODAS

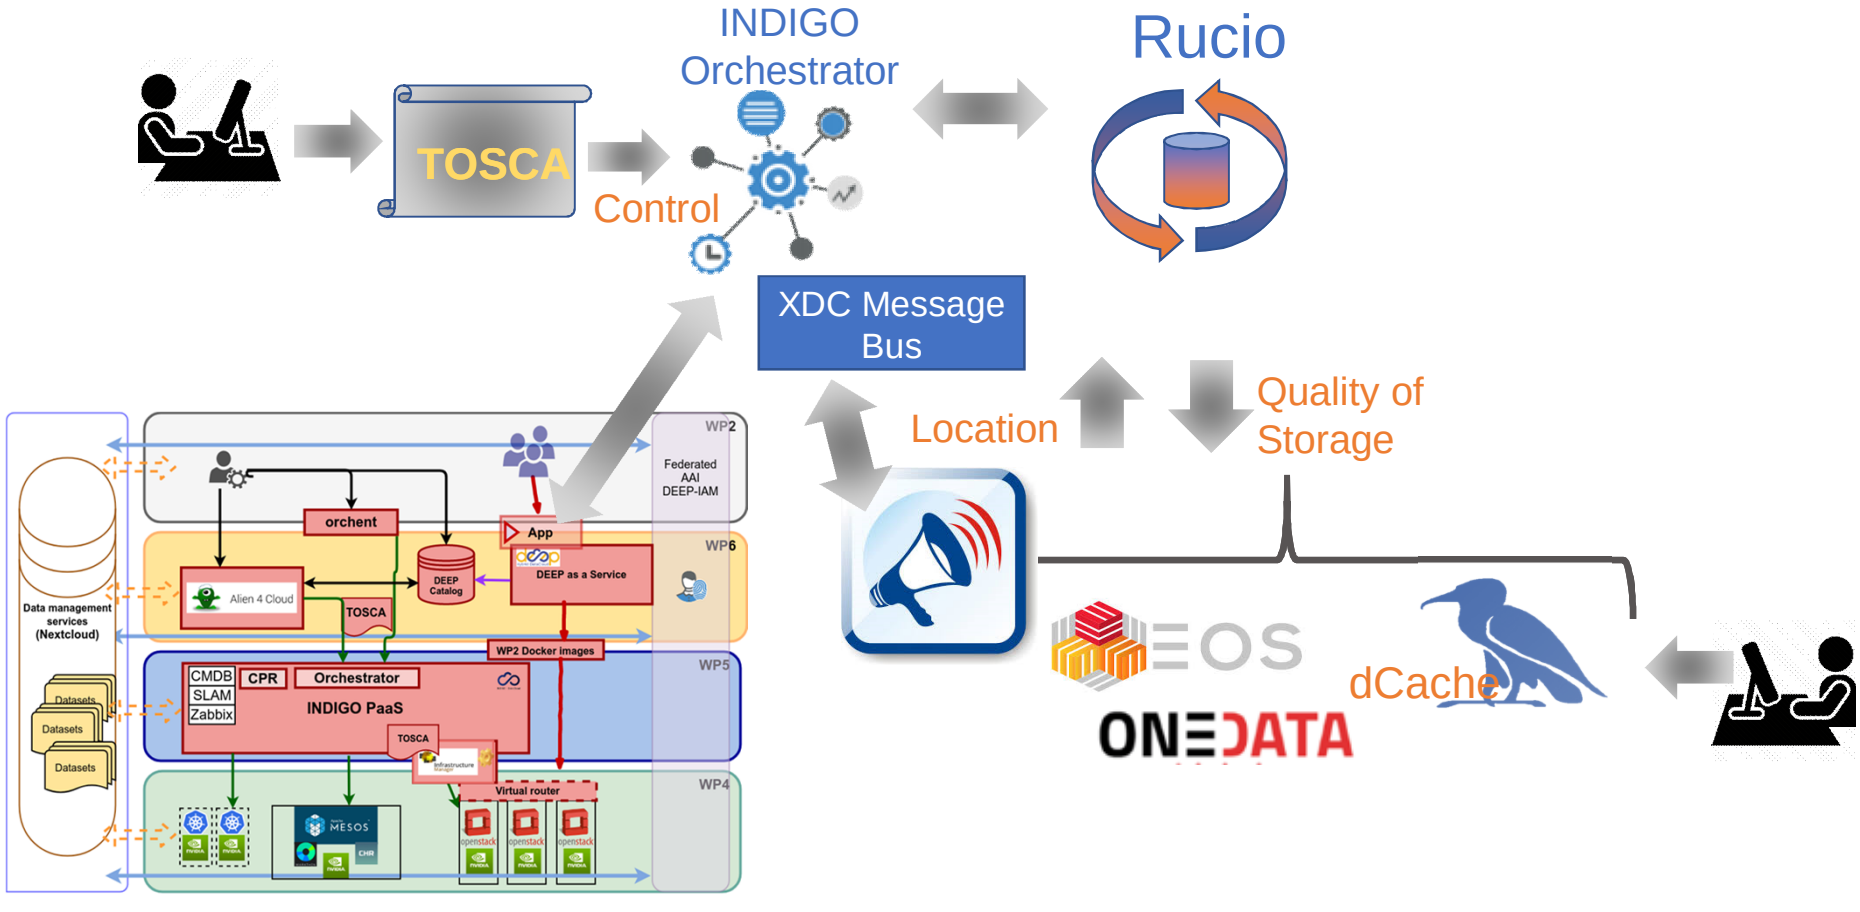

IoTwins

EOSC Ecosystem…the oversimplified story

EINFRA-21

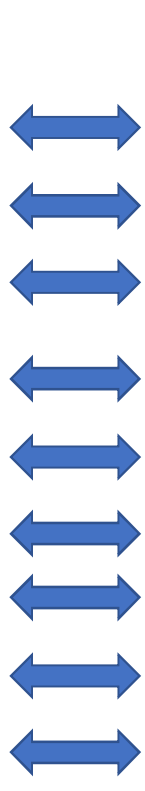

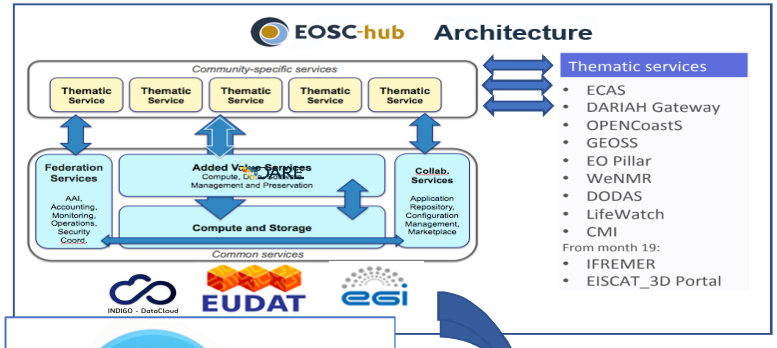

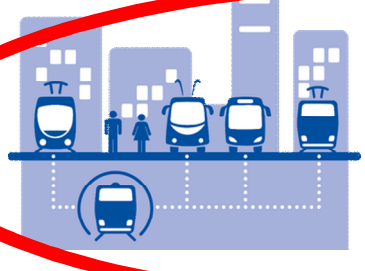

Communities/Experiments
E-infrastructures
Can or cannot be part of the EOSC
i.e. ESCAPE, HNSciCloud

Other infrastructures, not directly linked with the EOSC e-infra

IoTwins slides removed

Summary

Summary of the Use Cases vs project

Use Case	Addressed within	Used in project	Used by communities
Next gen. Identity Management	INDIGO-DC	XDC, DEEP, ESCAPE, IoTwins	WLCG
Base Cluster Inst. /HT-Condor-Batch Inst.	INDIGO-DC HUB	HUB, DEEP (ML)	CMS
HT-Condor-Batch + CVMFS	INDIGO-DC HUB	HUB	AMS, DAMPE VIRGO
Data federation management (xCache, EOS, dCache)	XDC	XDC, HNSciCloud, IoTwins, ESCAPE	WLCG XDC Communities
Spark+Hadoop cluster instantiation	INDIGO-DC HUB	DEEP, IoTwins	DEEP
Deep-Learning as a service	DEEP	IoTwins	DEEP communities
Policy driven data management based on QoS	XDC	IoTwins, ESCAPE, DEEP	All
Storage event notification handling	XDC	DEEP	Photon Science, DEEP Communities